



## DRAFT PCI-SIG ENGINEERING CHANGE NOTICE

<b>TITLE:</b>	Hierarchy ID Message
<b>DATE:</b>	Introduced: Feb 23, 2016 Updated: Sept 8, 2016 14-day Cross WG Review: Sept 15, 2016 Member Review: October 20, 2016 Final Approval: TBD
<b>AFFECTED DOCUMENT:</b>	PCI Express Base Specification, Version 3.1
<b>SPONSOR:</b>	Steve Glaser, Nvidia

**Part I****1. Summary of the Functional Changes**

Defines a new, optional PCISIG Defined Type 1 Vendor Defined Message.

This message provides software and/or firmware, running on a Function, additional information to uniquely identify that Function, within a large system or a collection of systems.

When a single system contains multiple PCI Express Hierarchies, this message tells a Function which Hierarchy it resides in. This value, in conjunction with the Routing ID number uniquely identifies a Function within that system.

In clustered system, this message can include a System Globally Unique Identifier (System GUID) for each system. This value, in conjunction with the Hierarchy ID and Routing ID uniquely identifies a Function within that cluster.

**2. Benefits as a Result of the Changes**

Currently, a Function knows its Routing ID, but it doesn't need to know which System or which Hierarchy within that System it belongs to. With the advent of clustered systems, and the potential for consolidating Sideband signals into a distinct bus (with its own addressing structure), it is helpful, in some systems, to provide additional information to help software determine the topology.

**3. Assessment of the Impact**

No impact to existing systems. No impact to future systems that don't need this functionality.

**4. Analysis of the Hardware Implications**

New optional message. New registers in config space.

**5. Analysis of the Software Implications**

No required changes. Software for clustered systems can benefit by using this mechanism. Future software that uses the consolidated Sideband signaling mechanism currently under discussion could also benefit.

## **6. Analysis of the C&I Test Implications**

Tests should be able to generate this message and see that supporting Endpoints record it. Optional Downstream Port tests should attempt to generate the message from a supporting Downstream Port and verify that the message was transmitted.

## Part II

### Detailed Description of the change

*In the PCI Express Base Specification, add new Section 6.x to Chapter 6 as follows:*

#### **6.x Hierarchy ID Message**

When software initializes a PCI Hierarchy, it assigns unique Bus and Device numbers to each component so that every Function in the Hierarchy has a unique Routing ID. Routing IDs are not unique in large Root Complexes that contain more than one Hierarchy and in clustered systems that contain multiple Root Complexes. Additional information is required to augment the Routing ID to produce a unique number. Functions can be uniquely identified by the combination of:

- ☐ Unique Identifier for the System (or Root Complex)
- ☐ Unique Identifier for the Hierarchy within that Root Complex
- ☐ Routing ID within that Hierarchy

The Hierarchy ID message is used to provide the additional information needed to uniquely identify a Function to itself.

Hierarchy ID Messages are sent by a Downstream Port upon software request. Received messages are reported in the Hierarchy ID Extended Capability (see Section 7.x).

Hierarchy ID Messages are Broadcast MsgD TLPs.

Hierarchy ID Messages are a PCISIG Defined Type 1 VDM. Hierarchy ID Messages can safely be sent at any time and components that do not comprehend them will silently ignore them.

Hierarchy ID Messages are typically sent from Root Ports. In systems where the Root Port does not support Hierarchy ID Messages, Hierarchy ID Messages can be sent from Switch Downstream Ports.

The Hierarchy ID Message is intended for use by software, firmware, and/or hardware. When using the Hierarchy ID Message, all bits of the Hierarchy ID, System GUID, System GUID Authority ID fields must be compared, without regard to any internal structure. How this information is used is outside the scope of this specification.

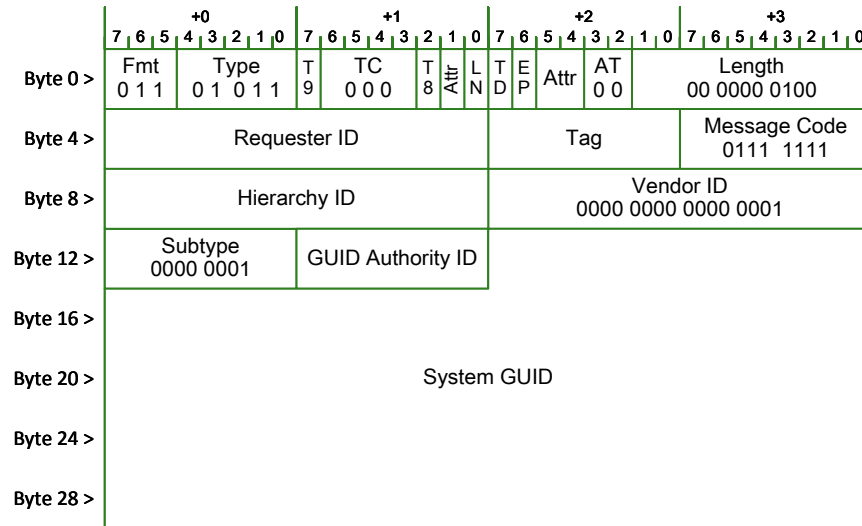


Figure 6-1: Hierarchy ID Message

**Hierarchy ID** contains the Segment Group Number associated with this Hierarchy (as defined by the *PCI Firmware Specification*). This field can be used in conjunction with the Routing ID to uniquely identify a Function within a System.

**System GUID**[143:0], in conjunction with System GUID Authority ID, provides a globally unique identification for a System.

System GUID[143:136] is byte 14 in the Hierarchy ID Message.  
 System GUID[135:128] is byte 15 in the Hierarchy ID Message.  
 System GUID[127:120] is byte 16 in the Hierarchy ID Message.  
 System GUID[119:112] is byte 17 in the Hierarchy ID Message.  
 System GUID[111:104] is byte 18 in the Hierarchy ID Message.  
 System GUID[103: 96] is byte 19 in the Hierarchy ID Message.  
 System GUID[ 95: 88] is byte 20 in the Hierarchy ID Message.  
 System GUID[ 87: 80] is byte 21 in the Hierarchy ID Message.  
 System GUID[ 79: 72] is byte 22 in the Hierarchy ID Message.  
 System GUID[ 71: 64] is byte 23 in the Hierarchy ID Message.  
 System GUID[ 63: 56] is byte 24 in the Hierarchy ID Message.  
 System GUID[ 55: 48] is byte 25 in the Hierarchy ID Message.  
 System GUID[ 47: 40] is byte 26 in the Hierarchy ID Message.  
 System GUID[ 39: 32] is byte 27 in the Hierarchy ID Message.  
 System GUID[ 31: 24] is byte 28 in the Hierarchy ID Message.  
 System GUID[ 23: 16] is byte 29 in the Hierarchy ID Message.  
 System GUID[ 15: 8] is byte 30 in the Hierarchy ID Message.  
 System GUID[ 7: 0] is byte 31 in the Hierarchy ID Message.

**System GUID Authority ID** identifies the mechanism used to ensure that the System GUID is globally unique. The mechanism for choosing which Authority ID to use for a given system is implementation specific. The defined values are shown in Table 6-1.

**Table 6-1: System GUID Authority ID Encoding**

Authority ID	Description
00h	<b>None</b> – System GUID[143:0] is not meaningful. System GUID[143:0] must be 0.
01h	<b>Timestamp</b> – System GUID[63:0] contains a timestamp associated with the particular system. Encoding is a Unix 64 bit time (number of seconds since midnight UTC January 1, 1970). The mechanism of choosing the timestamp to represent a system is implementation specific. System GUID[143:64] must be 0.
02h	<b>IEEE EUI-48</b> – System GUID[47:0] contains a 48 bit Extended Unique Identifier (EUI-48) associated with the particular system. Encoding is defined the the IEEE. See <a href="#">EUI-48<sup>1</sup></a> for details. EUI-48 values are frequently used as network interface MAC addresses. The mechanism of choosing the EUI-48 value to represent a system is implementation specific. System GUID[143:48] must be 0.
03h	<b>IEEE EUI-64</b> – System GUID[63:0] contains a 64 bit Extended Unique Identifier (EUI-64) associated with the particular system. Encoding is defined by the IEEE. See <a href="#">EUI-64<sup>2</sup></a> for details. The mechanism of choosing the EUI-64 value to represent a system is implementation specific. System GUID[143:64] must be 0.
04h	<b>RFC-4122 UUID</b> – System GUID[127:0] contain a UUID as defined by the IETF in <a href="#">RFC-4122<sup>3</sup></a> . This definition is technically equivalent to <a href="#">ITU-T Rec. X.667<sup>4</sup></a>   ISO/IEC 9834-8. The mechanism of choosing the UUID value to represent a system is implementation specific. System GUID[143:128] must be 0
05h	<b>IPv6 Address</b> – System GUID[127:0] contains the unique IPv6 address of one of the network interfaces of the system.

<sup>1</sup> <https://standards.ieee.org/develop/regauth/tut/eui48.pdf>

<sup>2</sup> <https://standards.ieee.org/develop/regauth/tut/eui64.pdf>

<sup>3</sup> <https://tools.ietf.org/html/rfc4122>

<sup>4</sup> <http://www.itu.int/rec/T-REC-X.667-201210-I/en>

	<p>The mechanism of choosing the IPv6 value to represent a system is implementation specific.</p> <p>System GUID[143:128] must be 0.</p>
06h to 7Fh	<p><b>Reserved</b> – System GUID[143:0] contains a unique value. The mechanism used ensure uniqueness is defined in a future version of this specification.</p>
80h to FFh	<p><b>PCISIG Vendor Specific</b> – System GUID Authority ID values 80h to FFh are reserved for PCISIG Vendor Specific usage.</p> <p>System GUID[143:128] contains a PCISIG assigned Vendor ID.</p> <p>System GUID[127:0] contain a unique number assigned by that vendor. The mechanism used for assigning numbers is implementation specific. One possible mechanism would be to use the serial number assigned to the system.</p> <p>The mechanism used to choose between these System GUID Authority IDs is implementation specific. One usage would be to allow a vendor to define up to 128 distinct 128-bit System GUID schemes.</p>



## IMPLEMENTATION NOTE

### System GUID Consistency and Stability

Software should ensure that a single system uses identical System GUID and System GUID Authority ID values everywhere. Doing otherwise defeats the purpose of System GUID.

Implementers should carefully consider their stability requirements for the System GUID value. For example, some use cases may require that the value not change when the system is rebooted. In those cases, a mechanism that picks the EUI-48 associated with first Ethernet MAC address discovered might be problematic if the result changes due to hardware failure, system reconfiguration, or variations/parallelism in the discovery algorithm.



## IMPLEMENTATION NOTE

### Hierarchy ID vs. Device Serial Number

The Device Serial Number mechanism can also be used to uniquely identify a component (see Section 7.17). Device Serial Number is a more expensive solution to this problem as it involves a ROM associated with each component.

*In the PCI Express Base Specification, add new Section 7.x to Chapter 7 as follows*

## 7.x Hierarchy ID Extended Capability

At most one instance of this capability is permitted in a Function.

## DRAFT-DRAFT-DRAFT-DRAFT-DRAFT-DRAFT-DRAFT-DRAFT

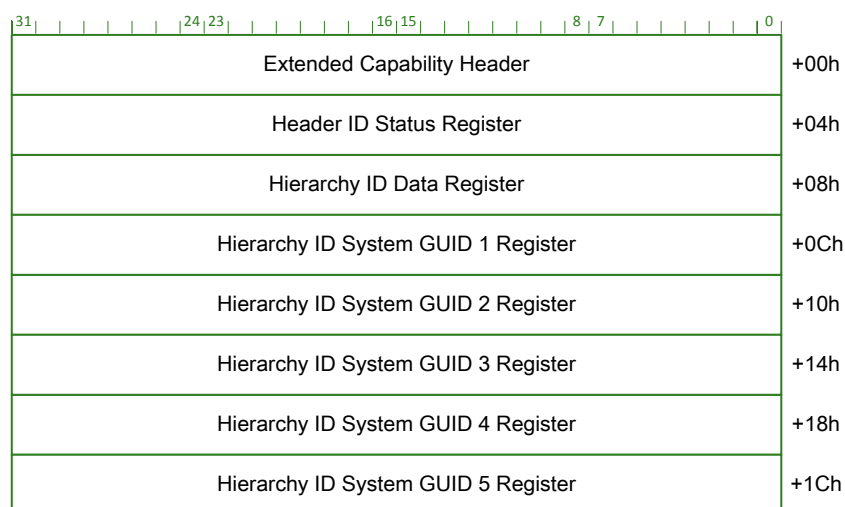
This capability takes two forms:

In Upstream Ports:

- ☐ This capability is permitted any Function associated with an Upstream Port.
- ☐ This capability is optional in Switch Upstream Ports. Support in Switch Upstream and Downstream Ports is independently optional.
- ☐ This capability is mandatory in Functions that use the Hierarchy ID Message. This includes use by the Function's driver.
- ☐ All Functions in an Upstream Port that implement this capability must report identical information.
- ☐ In Upstream Ports, the Downstream Port Routing ID, Hierarchy ID, System GUID Authority ID, and System GUID fields are Read Only and contain the values received in a Hierarchy ID Message.

In Downstream Ports:

- ☐ This capability is permitted in any Downstream Port. It is recommended that it be implemented in Root Ports.
- ☐ When present in a Switch Downstream Port, this capability must be implemented in all Downstream Ports of the Switch. Support in Switch Upstream and Downstream Ports is independently optional.
- ☐ In Downstream Ports, the Hierarchy ID, System GUID Authority ID, and System GUID fields are Read / Write and contain the values to send in the Hierarchy ID Message.
- ☐ A Hierarchy ID capability is not affected by Hierarchy ID Messages forwarded through the associated Downstream Port.



**Figure 7-1: Hierarchy ID Extended Capability**

## 7.x.1 Hierarchy ID Extended Capability Header




Figure 7-2: Hierarchy ID Extended Capability Header

Table 7-2: Hierarchy ID Extended Capability Header

Bit Location	Description	Attributes
15:0	<b>Extended Capability ID</b> – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability.  PCI Express Extended Capability ID for the Hierarchy ID Capability is 0028h..	RO
19:16	<b>Capability Version</b> – This field is a PCI-SIG defined version number that indicates the version of the Capability structure present.  Must be 1h for this version of the specification.	RO
31:20	<b>Next Capability Offset</b> – This field contains the offset to the next PCI Express Extended Capability structure or 000h if no other items exist in the linked list of Capabilities.	RO

## 7.x.2 Hierarchy ID Status Register




Figure 7-3: Hierarchy ID Status Register

Table 7-3: Hierarchy ID Status Register

Bit Location	Description	Attributes
15:0	<b>Downstream Port Routing ID</b> – In an Upstream Port, this field contains the Routing ID from the most recently received Hierarchy ID Message. This field is meaningful only if Hierarchy ID Valid is 1b.	RO/RsvdZ



	In a Downstream Port, this field is RsvdZ. This field defaults to 0000h.	
29:16	RsvdZ	RsvdZ
30	<b>Hierarchy ID Pending</b> – In a Downstream Port, this field is set to request that the Port send a Hierarchy ID Message. This bit is cleared when the message has been transmitted.  In an Upstream Port, this bit is RsvdZ. This bit defaults to 0b.	RW1S/RsvdZ
31	<b>Hierarchy ID Valid</b> – In an Upstream Port, this field is Set when a Hierarchy ID Message is received. It indicates that the remaining fields in this capability are meaningful.  In a Downstream Port, this bit is RsvdZ. This bit defaults to 0b.	RO/RsvdZ

### 7.x.3 Hierarchy ID Data Register




Figure 7-4: Hierarchy ID Data Register

Table 7-4: Hierarchy ID Data Register

Bit Location	Description	Attributes
7:0	<b>System GUID Authority ID</b> – This field contains the Authority ID qualifying the System GUID field in the Hierarchy ID Message. See Section 6.x for details.  In an Upstream Port, this field is RO and contains the Authority ID value from the most recently received Hierarchy ID Message. This field is meaningful only if Hierarchy ID Valid is 1b.  In a Downstream Port, this field is RW and contains the Authority ID value to be sent in the Hierarchy ID Message. This field defaults to 00h.	RO/RW
15:8	RsvdP	RsvdP

31:16	<p><b>Hierarchy ID</b> – This field contains a software created ID that uniquely identifies the PCI Express Hierarchy with the Root Complex. The value 0000h indicates the default Hierarchy of the Root Complex. Non-zero values indicate additional Hierarchies.</p> <p>In an Upstream Port, this field is RO and contains the Hierarchy ID value from the most recently received Hierarchy ID Message. This field is meaningful only if Hierarchy ID Valid is 1b.</p> <p>In a Downstream Port, this field is RW and contains the Hierarchy ID value to be sent in the Hierarchy ID Message. This field defaults to 0000h.</p>	RO/RW
-------	--	-------

## 7.x.4 Hierarchy ID GUID 1 Register




Figure 7-5: Hierarchy ID GUID 1 Register

Table 7-5: Hierarchy ID GUID 1 Register

Bit Location	Description	Attributes
15:0	<p><b>System GUID 1</b> – This field contains bits [143:128] of the unique ID associated with the Root Complex.</p> <p>In an Upstream Port, this field is RO, contains the System GUID bits from the most recently received Hierarchy ID Message, and is meaningful only if Hierarchy ID Valid is 1b.</p> <p>In a Downstream Port, this field is RW and contains the System GUID bits to be sent in the Hierarchy ID Message. This field defaults to 0000h.</p>	RO/RW
31:0	RsvdP	RsvdP

## 7.x.5 Hierarchy ID GUID 2 Register




Figure 7-6: Hierarchy ID GUID 2 Register

Table 7-6: Hierarchy ID GUID 2 Register

Bit Location	Description	Attributes
31:0	<p><b>System GUID 2</b> – This field contains bits [127:96] of the unique ID associated with the Root Complex.</p> <p>In an Upstream Port, this field is RO, contains the System GUID bits from the most recently received Hierarchy ID Message, and is meaningful only if Hierarchy ID Valid is 1b.</p> <p>In a Downstream Port, this field is RW and contains the System GUID bits to be sent in the Hierarchy ID Message.</p> <p>This field defaults to 0000 0000h.</p>	RO/RW

## 7.x.6 Hierarchy ID GUID 3 Register




Figure 7-7: Hierarchy ID GUID 3 Register

Table 7: Hierarchy ID GUID 3 Register

Bit Location	Description	Attributes
31:0	<p><b>System GUID 3</b> – This field contains bits [95:64] of the unique ID associated with the Root Complex.</p> <p>In an Upstream Port, this field is RO, contains the System GUID bits from the most recently received Hierarchy ID Message, and is meaningful only if Hierarchy ID Valid is 1b.</p> <p>In a Downstream Port, this field is RW and contains the System GUID bits to be sent in the Hierarchy ID Message.</p> <p>This field defaults to 0000 0000h.</p>	RO/RW

## 7.x.7 Hierarchy ID GUID 4 Register




Figure 7-8: Hierarchy ID GUID 4 Register

Table 7-8: Hierarchy ID GUID 4 Register

Bit Location	Description	Attributes
31:0	<p><b>System GUID 4</b> – This field contains bits [31:0] of the unique ID associated with the Root Complex.</p> <p>In an Upstream Port, this field is RO, contains the System GUID bits from the most recently received Hierarchy ID Message, and is meaningful only if Hierarchy ID Valid is 1b.</p> <p>In a Downstream Port, this field is RW and contains the System GUID bits to be sent in the Hierarchy ID Message.</p> <p>This field defaults to 0000 0000h.</p>	RO/RW

## 7.x.8 Hierarchy ID GUID 5 Register




Figure 7-9: Hierarchy ID GUID 5 Register

Table 7-9: Hierarchy ID GUID 5 Register

Bit Location	Description	Attributes
31:0	<p><b>System GUID 5</b> – This field contains bits [31:0] of the unique ID associated with the Root Complex.</p> <p>In an Upstream Port, this field is RO, contains the System GUID bits from the most recently received Hierarchy ID Message, and is meaningful only if Hierarchy ID Valid is 1b.</p> <p>In a Downstream Port, this field is RW and contains the System GUID bits to be sent in the Hierarchy ID Message.</p> <p>This field defaults to 0000 0000h.</p>	RO/RW

*In Appendix F, change Table F-2 as follows:*

**Table F-2: PCI-SIG-Defined VDM Subtype Usage**

Subtype	Routing r[2:0]	Type	Description
0000 0000	010 or 011	MsgD	LN Message, see Section 2.2.8.6.1.1
<u>0000 0001</u>	<u>011</u>	<u>MsgD</u>	<u>Hierarchy ID Message, See Section 6.x</u>
0111 1111	100	Msg	Device Readiness Status, see Section 2.2.8.6.3
0111 1111	000	Msg	Function Readiness Status, see Section 2.2.8.6.4